arXiv:1511.08303v1 [cs.DS] 26 Nov 2015

Engineering Oracles for Time-Dependent Road Networks*

Spyros Kontogiannis[†,◇] George Michalopoulos[‡,◇] Georgia Papastavrou[†,◇]
Andreas Paraskevopoulos[‡,◇] Dorothea Wagner[◊] Christos Zaroliagis[‡,◇]

Abstract

We implement and experimentally evaluate landmark-based oracles for min-cost paths in two different types of road networks with time-dependent arc-cost functions, based on distinct real-world historic traffic data: the road network for the metropolitan area of Berlin, and the national road network of Germany.

Our first contribution is a significant improvement on the implementation of the `FLAT` oracle, which was proposed and experimentally tested in previous works. Regarding the implementation, we exploit *parallelism* to reduce preprocessing time and real-time responsiveness to live-traffic reports. We also adopt a *lossless compression scheme* that severely reduces preprocessing space and time requirements. As for the experimentation, apart from employing the new data set of Germany, we also construct several *refinements* and *hybrids* of the most prominent landmark sets for the city of Berlin. A significant improvement to the speedup of `FLAT` is observed: For Berlin, the average query time can now be as small as 83μsec, achieving a speedup (against the time-dependent variant of Dijkstra's algorithm) of more than 1,119 in absolute running times and more than 1,570 in Dijkstra-ranks, with worst-case observed stretch less than 0.781%. For Germany, our experimental findings are analogous: The average query-response time can be 1.269msec, achieving a speedup of more than 902 in absolute running times, and 1,531 in Dijkstra-ranks, with worst-case stretch less than 1.534%.

Our second contribution is the implementation and experimental evaluation of a novel hierarchical oracle (`HORN`). It is based on a *hierarchy* of landmarks, with a few "global" landmarks at the top level possessing travel-time information for all possible destinations, and many more "local" landmarks at lower levels possessing travel-time information only for a small neighborhood of destinations around them. As it was previously proved, the advantage of `HORN` over `FLAT` is that it achieves query times sublinear, not just in the size of the network, but in the Dijkstra-rank of the query at hand, while requiring asymptotically similar preprocessing space and time. Our experimentation of `HORN` in Berlin indeed demonstrates improvements in query times (more than 30.37%), Dijkstra-ranks (more than 39.66%), and also worst-case error (more than 35.89%), at the expense of a small blow-up in space.

Finally, we implement and experimentally test a dynamic scheme to provide responsiveness to live-traffic reports of incidents with a small timelife (e.g., a temporary blockage of a road segment due to an accident). Our experiments also indicate that the traffic-related information can be updated in seconds.

1 Introduction

Min-cost-path oracles (a.k.a. distance oracles) are succinct data structures encoding min-cost path information among a carefully selected subset of pairs of vertices in a graph. The encoding is done in such a way that the oracle can efficiently answer min-cost path queries for arbitrary origin-destination pairs, exploiting the preprocessed data and/or local min-cost path searches. An oracle is exact (resp. approximate) if the returned costs by the accompanying query algorithm are exact (resp. approximate). A bulk of important work (e.g.,

*Partially supported by the EU FP7/2007-2013 under grant agreements no. 609026 (project MOVESMART) and no. 621133 (project HoPE), and by the DFG under grant agreements no. WA654/16-2 and no. WA654/23-1.

† Computer Science & Engineering Dept., University of Ioannina, 45110 Ioannina, GREECE. kontog@cs.uoi.gr, gioulycs@gmail.com.
‡ Computer Engineering & Informatics Dept., University of Patras, 26500 Rio, GREECE. {michalog,paraskevop,zaro}@ceid.upatras.gr.
◇ Computer Technology Institute and Press "Diophantus", 26504 Rio, GREECE.
◊ Karlsruhe Institute of Technology (KIT), Am Fasanengarten 5, 76131 Karlsruhe, GERMANY. dorothea.wagner@kit.edu

Last change: November 26, 2015. Compiled: November 30, 2015.

[22, 21, 16, 17, 23, 24, 1]) is devoted to constructing and analysing oracles for *static* (i.e., *time-independent*), mostly undirected networks in which the arc-costs are fixed scalars, providing trade-offs between the oracle's space and query time and, in case of approximate oracles, also of the approximation guarantee, typically tagged as the *stretch* (factor). For an overview of oracles for networks with static arc-cost metrics, the reader is deferred to [20].

Considerable experimental work on routing in large-scale road networks has also appeared in recent years, with remarkable achievements that have been demonstrated on continental-size road-network instances. The goal is again to preprocess the arc-cost metric and then propose query algorithms (known as *speedup techniques* in this framework) for responding to min-cost path queries in time that is several orders of magnitude faster than the execution of Dijkstra's algorithm. An excellent overview of this line of research is provided in [2]. Once more, the bulk of the literature concerns *static* metrics, with only a few exceptions [3, 4, 5, 14].

Nevertheless, contemporary transportation networks are typically considered as directed graphs accompanied by an arc-cost metric (e.g., travel-time, energy-consumption) which demonstrates *time-varying* behaviour. For example, critical road segments in the center, or at an entry point of a metropolitan area, suffer traversal times which are heavily dependent on the actual time of traversal. The rapid penetration of modern technology, such as the cellular technology, in our daily habits, has allowed travellers to act as moving sources of real-traffic information. Therefore, vehicle navigation and route planning vendors are nowadays in position to acquire in real time instantaneous live-traffic reports (e.g., road blockages due to car accidents), as well as periodic speed-probes that allow the maintenance of historic traffic data about this time-varying behaviour of each and every road segment in a road network. This latter kind of historic traffic data can be succinctly represented as *arc-traversal-time* functions, in particular the interpolants of the sampled arc-traversal-time values, for all the arcs in the network infrastructure. One grand challenge is then to assess the earliest-arrival time for a given origin-destination pair $(o, d) \in V \times V$ and departure time $t_o \geq 0$ from the origin. If the metric of arc-traversal-times possesses the so-called *FIFO* property, according to which a delayed departure from o cannot possibly lead to an earlier arrival at d, then all the classical shortest-path techniques (such as Dijkstra's and Bellman-Ford algorithms) have their time-dependent variants [7, 15]. Of course, these are not the choices one should consider as query algorithms, in case of a route planning service that has to reply in real-time to several dozens, or even hundreds, of queries within a large-scale network.

1.1 Related Work. Providing min-cost paths in graphs with time-dependent arc-costs is a hard problem. The FIFO property characterizes the complexity of the problem [15]. For FIFO-abiding networks, the first time-dependent variant of Dijkstra's algorithm (`TDD`) was proposed in [7].

A few time-dependent variants of well-known speedup techniques for road networks have appeared in the literature, with remarkable *empirical* performance in some cases. All of them were experimentally evaluated on benchmark instances of the time-dependent European and/or the real data set for the German road networks, kindly provided by PTV AG for scientific use. In particular, efficient approaches for responding to earliest-arrival-time queries are provided in [5, 6], by combining vertex contractions with goal-directed techniques. The currently best speedup technique for both *earliest-arrival-time* and *travel-time function* (a.k.a. *profile search*) computations in time-dependent road networks was given in [4]. Their average response times to arbitrary queries for the German road network, of $4.7M$ vertices and $10.8M$ arcs, is less than 1.5msec and preprocessing space requirements of less than $1GB$. A point-to-point travel-time function can be constructed in less than 40msec, when the departure times interval is a single day. For point-to-point approximate travel-time functions, with experimentally observed stretch at most 1%, the construction time is less than 3.2msec. Their approach is based on the *time-dependent Contraction Hierarchies* [3], along with several heuristic improvements on the preprocessing step and on the query method.

To our knowledge, oracles for time-dependent networks achieving *provable* approximation guarantees, small preprocessing-space complexity and sublinear query time complexity, have been only recently proposed and theoretically analysed [11, 10], and experimentally evaluated [9].

1.2 Our Contribution. We engineer and experimentally evaluate oracles for time-dependent road networks on two real instances corresponding to qualitatively different cases. The first one corresponds to a typical weekday of the metropolitan road network of Berlin (about $4.7K$ vertices and $1.13M$ arcs) and was kindly provided to us by TomTom [19] for scientific use. The second instance corresponds to a typical weekday of the German road

network (about $4.6M$ vertices and $11.18M$ arcs) and was kindly provided to us by [18], also for scientific use. More details on the identities of the two instances can be found in Table 2.

In this work, we make two main contributions. First, we significantly extended and improved the implementation and experimental evaluation of the FLAT oracle proposed in [11] and originally implemented and experimentally assessed in [9]. Regarding the implementation, we exploit *parallelism* to reduce preprocessing time and real-time responsiveness to live-traffic reports. We also adopt a *lossless compression scheme* that severely reduces preprocessing space and time requirements. As for the experimentation, apart from employing the new data set of Germany, we also construct several *refinements* and *hybrids* of the most prominent landmark sets (RANDOM and KAHIP) for the city of Berlin. A significant improvement to the speedup of FLAT is observed: the average query time can now be as small as 81μsec, if we exploit the query algorithm FCA along with the SPARSE-RANDOM landmark set, corresponding to a speedup against TDD that is larger than $1,146$ w.r.t. query-response times and more than $1,227$ in Dijkstra-ranks, with worst-case observed stretch less than 0.771%, for the Berlin instance. If we use the SPARSE-KAHIP landmark set, the absolute query time is slightly worse (83μsec) but the speedup w.r.t. the Dijkstra-ranks is now more than $1,570$, with worst-case error less than 0.781%.

Our experimental findings for Germany are analogous: the average response-time to arbitrary queries can be at most 1.269msec, achieving a speedup of more than 902 in absolute running times and $1,531$ in Dijkstra-ranks, with worst-case stretch at most 1.534%.

Our second contribution is the implementation and experimental evaluation of the novel HORN oracle [10], which is based on a *hierarchy* of landmarks, from a few "global" landmarks possessing knowledge of the entire network towards (many more) "local" landmarks whose knowledge of the network is restricted to small neighborhoods around them. As was proved in [10], the advantage of HORN over FLAT is that it achieves query times *sublinear*, not just in the size of the network, but in the actual Dijkstra-rank of the query at hand, be it long-range, mid-range, or short-range, while requiring similar preprocessing space and time. Our experiments on the Berlin instance indeed confirm the improved stretch factors, but also better speedups due to sophisticated early-stopping criteria, compared to the experimentation on FLAT for the same subsets of "global" landmarks. In particular, our experimentation of HORN in Berlin indeed demonstrates improvements in query times (more than 30.37%), Dijkstra-ranks (more than 39.66%), and also worst-case error (more than 35.89%), at the expense of a small blow-up in space.

An additional contribution concerns the implementation and experimental evaluation of a dynamic scheme to provide responsiveness to live-traffic reports of incidents with a small timelife (e.g., a temporary blockage of a road segment due to an accident). Our experiments also indicated that traffic-related information can be updated in seconds.

2 Preliminaries

2.1 Notation. We consider directed graphs $G = (V, A)$ with $|V| = n$ vertices and $|A| = m = \mathcal{O}(n)$ arcs, where each arc $a \in A$ is accompanied with a continuous, periodic, piecewise linear (pwl) *arc-travel-time* (or *arc-delay*) function defined as follows: $\forall k \in \mathbb{N}, \forall t \in [0, T),\ D[a](kT + t) = d[a](t)$, where $d[a] : [0, T) \to [1, M_a]$ such that $\lim_{t\uparrow T} d[a](t) = d[a](0)$, for some fixed integer M_a denoting the maximum possible travel time ever observed at arc a. Notice that the minimum arc travel time value in the entire network is also normalized to 1. Each arc-travel-time function $D[a]$ can be represented succinctly as a list of K_a breakpoints defining $d[a]$. Let $K = \sum_{a\in A} K_a$ be the number of breakpoints to represent all of them, $K_{\max} = \max_{a\in A} K_a$, and K^* be the number of *concavity-spoiling* breakpoints, i.e., those in which the arc-travel-time slopes increase. Clearly, $K^* \leq K$, and $K^* = 0$ for *concave* pwl arc-travel-time functions.

The *arc-arrival-time* functions are defined as $Arr[a](t) = t + D[a](t),\ \forall t \in [0, \infty)$. A typical assumption is that each arc-arrival-time functions are strictly increasing, in order to satisfy the strict FIFO property. The *path-arrival-time* function of a path $p = \langle a_1, \ldots, a_k\rangle$ in G (represented as a sequence of arcs) is defined as the composition of the arc-arrival-time functions for the constituent arcs of p: $Arr[p](t) = Arr[a_k](Arr[a_{k-1}](\cdots(Arr[a_1](t))\cdots))$. The *path-travel-time* function is $D[p](t) = Arr[p](t) - t$. Also, between any origin-destination $(o, d) \in V \times V$, $\mathcal{P}_{o,d}$ is the set of all $od-$paths in G, and the *earliest-arrival-time / shortest-travel-time* functions are defined as follows: $\forall t_o \geq 0$, $Arr[o, d](t_o) = \min_{p\in\mathcal{P}_{o,d}}\{Arr[p](t_o)\}$ and $D[o, d](t_o) = \min_{p\in\mathcal{P}_{o,d}}\{D[p](t_o)\} = Arr[o, d](t_o) - t_o$. The *Dijkstra-rank* $\Gamma[o, d](t_o)$ is the number of settled vertices up to d, when executing the time-dependent variant of Dijkstra (we call it TDD) from (o, t_o). $SP[o, d](t_o) = \{p \in P_{o,d} : Arr[p](t_o) = Arr[o, d](t_o)\}$ is the set of minimum-travel-time paths for the triple

(o, d, t_o). $ASP[o,d](t_o)$ is the set of *od*-paths which are $(1+\varepsilon)$-approximations of minimum-travel-time *od*-paths.

For any $a = uv \in A$ and subinterval $[t_s, t_f) \subseteq [0, T)$, we define upper-bounding and lower-bounding (static) travel-time metrics: the *free-flow* travel-time $\underline{D}[uv](t_s, t_f) := \min_{t_u \in [t_s, t_f)}\{D[uv](t_u)\}$ and the *maximally-congested* travel-time $\overline{D}[uv](t_s, t_f) := \max_{t_u \in [t_s, t_f)}\{D[uv](t_u)\}$. If $[t_s, t_f) = [0, T)$, we refer to the static *free-flow* and *full-congestion* metrics $\underline{D}, \overline{D} : A \to [1, M]$, respectively. With respect to any arc-cost metric D, $diam(G, D)$ is the diameter of the graph, i.e., the largest possible distance in the graph. For example, $diam(G, \underline{D})$ is the free-flow diameter of G.

For any vertex $v \in V$, departure-time $t_v \in \mathbb{R}_{\geq 0}$ and integer $F \in [n]$, $B[v; F](t_v)$ ($B[v; R](t_v)$) is a ball of size $F \in \mathbb{N}$ (of radius $R > 0$) growing by TDD from (v, t_v) in the time-dependent metric. Analogously, $\underline{B}[v; F]$ ($\underline{B}[v; R]$), and $\overline{B}[v; F]$ ($\overline{B}[v; R]$) are the size-F (radius-R) balls from v under the (static) free-flow and fully-congested travel-time metrics.

A $(1+\varepsilon)$-*upper-approximation* $\overline{\Delta}[o, d]$ and a $(1+\varepsilon)$-*lower-approximation* $\underline{\Delta}[o, d]$ of $D[o, d]$, are continuous, pwl, periodic functions, with a (hopefully small) number of breakpoints in $[0, T)$, such that: $\forall t_o \geq 0,\ D[o,d](t_o) \ / \ (1+\varepsilon) \leq \underline{\Delta}[o,d](t_o) \leq D[o,d](t_o) \leq \overline{\Delta}[o,d](t_o) \leq (1+\varepsilon) \cdot D[o,d](t_o)$.

For convenience, the notation used throughout the paper is summarized in Table 1.

2.2 Assumptions on Travel-Time Metric. We adopt two assumptions from [11] and an additional one from [10] on the kind of shortest-travel-time functions that may appear in the time-dependent network instance at hand. All of them are quite natural and justified in time-depenent road networks. Technically, these assumptions allow the *smooth transition* from static metrics on undirected graphs towards time-dependent metrics on directed graphs. The first assumption asserts that all the minimum-travel-time slopes are bounded in a given interval $[-\Lambda_{\min}, \Lambda_{\max}]$.

ASSUMPTION 2.1. $\exists \Lambda_{\min} \in [0, 1),\ \exists \Lambda_{\max} \geq 0$ *s.t. the following holds:* $\forall (o, d) \in V \times V,\ \forall 0 \leq t_1 < t_2,\ (D[o,d](t_1) - D[o,d](t_2)) \ / \ (t_1 - t_2) \in [-\Lambda_{\min}, \Lambda_{\max}]$.

The lower-bound of -1 in the minimum-travel-time function slopes is indeed a direct consequence of the FIFO property, which is typically assumed to hold in several time-dependent networks, such as road networks. $\Lambda_{\max}$ represents the maximum possible rate of change of minimum-travel-times in the network, which only makes sense to be bounded (in particular, independent of the network size) in realistic instances such as the ones representing urban-traffic time-dependent road networks.

The next assumption states that the ratio of minimum-travel-times in opposite directions between two vertices, for any departure-time but not necessarily via the same path, is bounded by a constant.

ASSUMPTION 2.2. $\exists \zeta \geq 1,\ \forall (o, d) \in V \times V,\ \forall t \in [0, T),\ D[o,d](t) \leq \zeta \cdot D[d,o](t)$.

The last assumption states that the free-flow ball from a vertex blows-up by at most a polylogarithmic factor, if it is expanded up to the full-congestion radius within it, under the free-flow metric.

ASSUMPTION 2.3. *For any vertex $\ell \in V$, and a positive integer F, consider the (static) Dijkstra ball $\underline{B}[\ell; F]$ around ℓ under the free-flow metric. Let $\underline{R}[\ell] = \max\{\underline{D}[\ell, v] : v \in \underline{B}[\ell; F]\}$ and $\overline{R}[\ell] = \max\{\overline{D}[\ell, v] : v \in \underline{B}[\ell; F]\}$ be the largest free-flow and full-congestion travel-times from ℓ to any other vertex in $\underline{B}[\ell; F]$. Finally, let $\underline{B}'[\ell; F] = \{v \in V : \underline{D}[\ell, v](0, T) \leq \overline{R}[\ell]\}$ be the free-flow ball around ℓ with the (larger) radius $\overline{R}[\ell]$. Then it holds that $|\underline{B}'[\ell; F]| \in \mathcal{O}(F \cdot \text{polylog}(F))$.*

We conducted extended experimental analysis for two distinct road networks with time-dependent travel-times, the metropolitan area of Berlin, and the national road network of Germany. The former was kindly provided to us by TomTom [19] and the latter by PTV AG [18], for scientific purposes. More details on the identities of our instances can be found in Table 2. We verified the validity of these assumptions (cf. Tables 2, 3, and 4). The maximum observed values for ζ, $\Lambda_{\max}$ and the blow-up factor in free-flow ball size are 1.189, 0.216 and 8.3, respectively. More details are provided in Section 6.2.

3 Time-Dependent Oracles

3.1 The FLAT Oracle. The main rationale of the FLAT oracle [10] is the following: we somehow select a small subset of landmark vertices L. Then, we compute *travel-time summaries*, i.e., $(1+\varepsilon)-$approximations of

minimum-travel-time *functions*, from landmarks towards all reachable vertices. Finally, we use properly designed query algorithms that exploit these travel-time summaries in order to provide approximately minimum-travel-time *values* for arbitrary queries $(o, d, t_o) \in V \times V \times [0, T)$.

3.1.1 Selection of Landmark Sets. FLAT was preliminarily implemented and evaluated in [9] for six landmark sets of the Berlin instance. We considered three selection methods (RANDOM, METIS [13] and KAHIP [8]) and two sizes (1,000 and 2,000 landmarks). The experimental evaluation in [9] indicated that the RANDOM (R) and KAHIP (K) landmark sets dominated the rest, the former w.r.t. the query times and the latter w.r.t. observed relative errors. In this work we choose to focus on these two prominent landmark sets.

For the Berlin instance, we create landmark sets, each containing 2,000 landmarks. The space requirements were originally quite demanding. After the lossless compression techniques that we adopt here (cf. Section 4) we managed to decrease the space by 2/3. We consider also some refinements of the R- and K-landmark sets:

- IMPORTANT-RANDOM (IR) landmarks: We take as input the R-landmark set and from each landmark we grow a small ball in the free-flow metric. Within this ball we move the landmark from the ball center to a vertex adjacent to the most significant road segment in the ball.
- SPARSE-RANDOM (SR) landmarks: We choose the landmarks again randomly but sequentially: Each time a landmark is selected, a small ball is subtracted from the network prior to the selection of the next landmark, so as to assure that the landmarks are well separated.
- SPARSE-KAHIP (SK) landmarks: We initially choose a superset of candidate-landmarks from a fine-grained KaHIP partition. The final landmark set is chosen from these candidates, once more randomly and sequentially: Each time a landmark is selected, a small ball of candidate-landmarks around it is also subtracted prior to the selection of the next landmark, so as to assure that the landmarks are well separated.
- HYBRID (H) landmarks: This landmark selection method combines the rationales of RANDOM and KAHIP. We start with a KAHIP partition that provides roughly 1,000 boundary vertices as landmarks. Another collection of roughly 1,000 additional landmarks is selected, in a balanced way among cells and uniformly at random within each cell of the KAHIP partition.

Finally, for the sake of the HORN oracle, we also construct hierarchical landmark sets, HIERARCHICAL-RANDOM (HR) and HIERARCHICAL-SPARSE-RANDOM (HSR) of 10,256 and 20,513 landmarks.

For the instance of Germany we only consider the R-, K-, SR- and SK-landmark sets, each of roughly 2,000 landmarks. Unfortunately we could not construct an IR-landmark set to this instance, since there is no such information concerning the classification of road segments.

3.1.2 Preprocessing Phase. After L is determined, a preprocessing phase is performed in which all $(1+\varepsilon)-$upper-approximating travel-time functions (travel-time summaries) from landmarks $\ell \in L$ towards destinations $\forall v \in V$ are computed and stored, based solely on the TRAP method [10]. TRAP splits the entire period $[0, T)$ into small, consecutive subintervals, each of length $\tau > 0$, and provides a crude approximation of the unknown shortest-travel-time function in each interval, based solely on Assumption 2.1. After sampling the travel-time values of each destination $v \in V$, for a given origin $u \in V$, we consider each pair of consecutive sampling times $t_s < t_f$ and the semilines with slopes $\Lambda_{\max}$ from t_s and $-\Lambda_{\min}$ from t_f. The considered upper-approximating function $\overline{D}[u, v]$ within $[t_s, t_f)$ is then (a refinement of) the lower-envelope of these two lines.

3.1.3 Query Algorithms. We consider the three query algorithms FCA, RQA, and FCA$^+(N)$. The first two were introduced in [11], while the third one was introduced in [9]. All algorithms can be fine-tuned to run in $o(n)$ time. A preliminary implementation and experimental evaluation was provided in [9].

FCA grows a ball $B_o \equiv B[o](t_o) = \{x \in V : D[o, x](t_o) \leq D[o, \ell_o](t_o)\}$ from (o, t_o), by running TDD until either d or the closest landmark $\ell_o \in \arg\min_{\ell \in L}\{D[o, \ell](t_o)\}$ is settled. It then returns either the exact travel-time value, or an $(1+\varepsilon+\psi)$-approximate travel-time value via ℓ_o, where ψ is a constant depending on ε, ζ, and $\Lambda_{\max}$, but not on n.

FCA$^+(N)$ is a variant of FCA that keeps growing a TDD ball from (o, t_o) until either d, or a given number N of landmarks is settled, and then returns the smallest via-landmark approximate travel-time value (among all N

settled landmarks). The approximation guarantee is the same as that of FCA, but its practical performance is impressive (in most cases even better than RQA).

RQA is a PTAS, providing an approximation guarantee of $1+\sigma = 1+\varepsilon \cdot \left[(1+\varepsilon/\psi)^{r+1}\right] / \left[(1+\varepsilon/\psi)^{r+1}-1\right]$, by exploiting carefully a number $r \in \mathbb{N}$ (called the *recursion budget*) of recursive accesses to the preprocessed information, each of which produces (via calls to FCA) additional candidate $od-$paths sol_i. RQA works as follows. As long as the destination vertex within the explored area around the origin has not yet been discovered, and there is still some remaining recursion budget, it "guesses" (by exhaustively searching for it) the next vertex w_k of the boundary set of touched vertices (i.e., still in the priority queue) along the unknown shortest $od-$path. Then it grows an outgrowing TDD ball from the new center $(w_k, t_k = t_o + D[o, w_k](t_o))$, until it reaches the closest landmark ℓ_k to it, at travel-time $R_k = D[w_k, \ell_k](t_k)$. This new landmark offers an alternative $od-$path sol_k by a new application of FCA.

In [9] we executed $\text{FCA}^+(N)$ *after* the execution of RQA, so as to set the value of N to the actual number of landmarks that were discovered by RQA, in each query. To avoid this privileged treatment of $\text{FCA}^+(N)$, we choose here to test it for several but fixed values of N, regardless of RQA's behavior.

3.2 The HORN Oracle. The novelty of the HORN oracle [10] is to create a hierarchy of landmark sets, whose range of "preprocessed destinations" gradually ranges from a few "nearby" vertices up to all reachable vertices (in the last level), in order to serve each query (o, d, t_o) only with relevant landmarks with respect to its own Dijkstra-rank $\Gamma[o, d](t_o)$. This way, we aim at achieving speedups similar to those of FLAT for long-range queries, to all possible ranges of queries, while increasing the space requirements only by a small factor. Our goal is to "guess" the order of the Dijkstra Rank $\Gamma[o, d](t_o)$ for (o, d, t_o). The guessing is achieved in a way that is typical in online algorithms that have to deal with an unknown parameter: Starting from a small value (say, $\mathcal{O}(\sqrt{n})$), we keep growing a ball from (o, t_o), increasing appropriately the value of the guess as the ball grows, until the very first time at which a successful completion of a proper variant of RQA is very likely to occur (exactly because we "guessed right" the actual Dijkstra-rank). The travel-time returned is that of the best possible *od*-path among all the successfully discovered approximate *od*-paths so far, via "informed" landmarks that possess travel-time summaries for d. The crux is in organizing the preprocessed information in such a way that it is indeed possible for the query algorithm to successfully complete its execution as soon as the "guess" asymptotically matches the value of $\Gamma[o, d](t_o)$.

The *Hierarchical Query Algorithm* (HQA) for (o, d, t_o) proceeds as follows: a single ball grows from (o, t_o), until either d is reached, or an *Early Stopping criterion* (ESC) is fulfilled, or the *Appropriate Level of Hierarchy* (ALH) of landmarks is reached (whichever occurs first). If d is settled by the ball from (o, t_o), an exact solution is returned. If ESC causes HQA to terminate, then the value $D[o, \ell_o](t_o) + \overline{\Delta}[\ell_o, d](t_o + D[o, \ell_o](t_o))$ is reported, because it is already a very good approximation. Otherwise, HQA, due to ALH, considers being at the right level-i of the hierarchy and continues executing the corresponding variant of RQA, call it RQA_i, which uses as its own landmark set $M_i = \cup_{j=i}^{4} L_j$. Observe that RQA_i may now fail constructing approximate shortest paths via certain landmarks in M_i that it settles, since they may not possess a travel-time summary for d. HQA terminates by returning the best *od*-path that has been discovered so far, via *all* settled landmarks which are "informed" (i.e., they have d in their coverage), either by the very first ball from (o, t_o) or by RQA_i. HQA uses some parameters: a is the degree of sublinearity in the query time, compared to the targeted Dijkstra-rank; β is related to the approximation guarantee achieved upon exit due to ESC; γ has to do with the number of levels that we create in the hierarchy; and ξ is the amount of slackness that we introduce in the size of the area of coverage. We set these parameters here to the values $a = 1,\ \beta = 1,\ \gamma = 1.88,\ \xi = 0.1$. A more detailed explanation of HQA, as well as of its parameters, is provided in [10].

4 Compressing Preprocessing Space

Due to the criticality of the preprocessing space, our goal is to achieve an efficient storage of the constructed travel-time summaries, while keeping a sufficient precision. The key is that some specific features can be exploited in order to reduce the required space.

4.1 Fixed Range. For a one-day time period, departure-times and arrival-times have a bounded value range. The same also holds for travel times which are at most one-day for any query within a country area such as Germany. Therefore, when the considered precision of the traffic data is within seconds, we handle time-values as

integers in the range $[0\ ,\ 86,399]$, for milliseconds as integers in $[0\ ,\ 86,399,999]$, etc. Any (real) time value within a single-day period, represented as a floating-point number t_f, can thus be converted to an integer t_i with fewer bytes and a given unit of measure. For a unit measure (or scale factor) s, the resulting integer is $t_i = \lceil t_f/s \rceil$. In this manner, t_i needs size $\lceil \log_2(t_f/s)/8 \rceil$ bytes. The division t_f/s has quotient π and remainder υ. Thus, $t_f = s \cdot \pi + \upsilon$ and $t_i = \lceil (s \cdot \pi + \upsilon)/s \rceil = \lceil \pi + \upsilon/s \rceil$, with $\upsilon < s$. Therefore, converting t_f to t_i results to an absolute error of at most $2s$. In the reverse process, for extracting the stored value, the conversion is $t_f^{'} = t_i \cdot s$.

4.2 Bucketing. The number of breakpoints of the arc-travel-time functions is a major factor of space increase on the resulting minimum-travel-time functions. A way to deal with this is by merging consecutive breakpoints having absolute difference in travel-time values less than few seconds, in each arc travel time function. In this manner, to preserve the upper bound error, each resulted breakpoint gets the largest travel time among the breakpoints which take part in merge. Depending on the bucketing parameter c, we can decrease the number of breakpoints, sacrificing part of accuracy. In our experiments, the bucketing led to the highest reduction (about 86%) in space requirements.

4.3 Piecewise Composition. Many shortest *od*-paths typicaly contain at least one arc with pwl travel-time function, making $D[o,d](t)$ also a pwl function. To avoid the space increase from storing breakpoints unrestrictedly, we analyse any such shortest *od*-path into two subpaths *o-p-d*. *p* is selected so that the *pd*-subpath is the maximum subpath with no arc having pwl travel-time function. Such *pd*-subpaths exist because the number of constant arc-travel-time functions is much larger than the number of the pwl ones. Thus, for $D[o,d](t)$ we only store a "predecessor" pointer to $D[o,p](t)$ and the *constant* travel-time offset $D[p,d]$ i.e., $D[o,d](t) = D[o,p](t) + D[p,d]$. In our experiments, this method lead to 40% reduction of the space requirements.

4.4 Delay Shifts. There are many shortest *od*-paths with travel-time $D[o,d](t)$ with delay variation. Thus we further reduced the required space as follows. The delay fluctuates around a constant value. By taking the minimum delay value, the leg-delays can be represented as small shifts from this value. Those small shifts, belonging to a smaller value range, can be stored even in 1 byte. This conversion led to more than 5% reduction of space requirements.

4.5 Compression. Since there is no need for all landmarks to be concurrently active, we can compress their data blocks. We used the library *zlib* for this compression, which led to 10% reduction in space.

4.6 Required space. For Berlin, the required space was limited to an average size of less than 14MB per landmark. For storing the time-values of approximate travel-time summaries, we considered 2.64sec as resolution, corresponding to a scale factor $s = 1.32$ (when counting time in seconds), which requires 2 bytes per time-value. For Germany, the required space was limited to average size of 25.7MB per landmark. For storing the time-values of approximate travel-time summaries, we considered 17.64sec as resolution, corresponding to a scale factor $s = 1.32$, and a bucketing factor of $c = 15$sec.

4.7 Indexing Travel-Time Summaries. For retrieving efficiently the required minimum travel-time function $D[\ell,d](t)$ from a landmark ℓ to a destination-node d, we need also to store an index. Depending on the oracle, we used two types of indices.

Flat Index. We maintain a vector of pointers per landmark, one pointer equals per destination. The pointer of destination v provides the address of the $D[\ell,v](t)$ data. The pointers are in ascending order of node ID. The search time is $\mathcal{O}(1)$ and the required space is $\mathcal{O}(n|L|)$.

Horn Index. For each node v we maintain a vector of pointers. The number of pointers equals to the number of landmarks, from any partition level, which have travel time to the corresponding node v as destination. Obviously, there exist at least one landmark from the highest partition level. The pointer of the associated landmark ℓ provides the address of the $D[\ell,v](t)$ data. The pointers are stored in ascending order of node ID. The search time is $\mathcal{O}(\log(|L|))$ and the required space is $\mathcal{O}(n|L|)$.

5 Live Traffic Reporting

In a server-side routing service that responds to several queries in real-time, various disruptions may occur "on the fly" (e.g., the abrupt and unforeseen congestion, or even blockage of a road segment for half an hour due to a car accident) and have to be taken into account for the affected route plans that have already been suggested or will be suggested in the near future. We thus consider dynamic scenarios where there is a stream of live-traffic reports about abnormal delays on certain road segments (arcs), along with a time-window $[r_s, r_e]$, of typically small duration, in which the disruption occurs.

Our update step involves the recomputation of travel-time summaries for a subset of landmarks in the vicinity of the disruption. In particular, for a disrupted arc $a = uv$ of disruption duration $[r_s, r_e]$, we run a (static) `Backward-Dijkstra` from u under the free-flow metric, with travel time radius of at most $r_e - r_s$. The limited travel time radius is used to trace only the nearest landmarks that may actually be affected by the disruption, leaving unaffected all the "faraway" landmarks. The goal is to update as soon as possible the recommendations for the drivers who are close to the area of disruption. For each affected landmark ℓ, we consider a *disruption-times window* $[t_s, t_e]$, containing the latest departure times from ℓ for arriving at the tail u at any time in the interval $[r_s, r_e]$ in which the disruption occurs. We then compute *temporal* travel-time summaries for each affected landmark and disruption-times window. This computation is conducted as in the preprocessing phase. Using a 15-min radius for the disruptions, we executed $1,000$ live-traffic updates for the instances of Berlin and Germany, each with $2,000$ SR-landmarks. Each disruption has duration at most 2h. The average number of affected landmarks was 32 for Berlin, and only 3 for Germany (mainly due to sparsity). The corresponding update times for their preprocessed data were 21sec and 37sec, respectively.

6 Experimental Evaluation

6.1 Experimental Setup. All algorithms were implemented using C++ (gcc version 4.8.2). To support all graph-operations we used the PGL library [12]. All experiments were executed by a Intel(R) Xeon(R) CPU E5-2643v3 3.40GHz using 128GB of RAM, on Ubuntu 14.04 LTS. We used 6 threads for the parallelism of the preprocessing phase. The query algorithms were executed on a single thread.

6.2 Identity of Instances. Table 2 demonstrates the statistics of the arc-travel-time functions in the two instances.

Table 3 demonstrates some statistics of the minimum-travel-times related to Assumptions 2.1 and 2.2. These statistics are deduced as follows: We run $10,000$ random queries in departure range R_{dep}=[9:00-20:00] (densely for rush hours 9:00-12:00). Each query consists of a random pair of nodes (o, d) and a random departure time $t \in R_{dep}$. We compute $D[o,d](t)$ and $D[d,o](t)$ and we set $\zeta[o,d] = \max\{D[o,d](t), D[d,o](t)\} \ / \ \min\{D[o,d](t), D[d,o](t)\}$. $\zeta_{avg} = \sum_{(o,d,t)} \zeta[o,d](t)$ is the average value of ζ and $\zeta_{\max} = \max_{(o,d,t)}\{\zeta[o,d](t)\}$ the maximum value of ζ, over all these random queries. We also present the maximum and minimum observed values of the travel-time slopes $\Lambda_{\max}$ and $-\Lambda_{\min}$ in functions $D[o,d]$, as well as the offset of the computed minimum-travel-time functions.

We now discuss the validity of Assumption 2.3. Table 4 summarizes the observed value of the free-flow blowup that we found in the two instances.

We compute these statistics as folllows: For a set of $5,000$ randomly chosen origins, we grow minimum-travel-time balls by running (static) `Dijkstra` in the *free-flow* metric, until a number F of destinations is settled around each origin. The first column denoted by $FF\ rank$ concerns exactly the different values for F that we consider as the size for these balls. In each of the size-F free-flow balls we compute the maximum min-travel-time t_{CG} of a leaf from its origin, in the *full-congestion* metric this time. We then extend each of the previous balls (again in the *free-flow* metric), until the free-flow radius reaches exactly the value t_{CG}. The size of the resulting free-flow ball is denoted by $ExtFF\ Rank$. We provide both the average values avg($ExtFF\ Rank$) and the maximum values max($ExtFF\ Rank$) observed for this parameter, in the instances of Berlin and Germany. As indicated in Table 4, the ratio between max($ExtFF\ Rank$) and $FF\ Rank$ is some constant never exceeding the value 8.3. The ratio when considering average values never exceeds the value 2.

6.3 Experimental Evaluation. We report here the outcome of our experiments on the instances of Berlin and Germany. The instance of Berlin consists of $473,253$ nodes and $1,126,468$ arcs. We conducted a preprocessing of the instance that contracts nodes which are not junctions (i.e., of degree 2 in the undirected graph). We also added some shortcut edges at the endpoints of chains of contracted nodes. This led to an amount of $292,356$

active nodes and 752,362 active edges in the instance. As for the instance of Germany, it consists of 4,692,091 nodes and 11,183,060 arcs. After the contraction of degree-2 nodes, we got an instance with 3,431,213 active nodes and 8,554,840 active edges.

6.3.1 `FLAT` **@ Berlin.** Tables 5 and 6 summarize the performance of the basic query algorithms of `FLAT` with respect to absolute running times and Dijkstra-rank values, respectively, for landmark sets of various sizes. The best performance per query algorithm is indicated by highlighted table cells. The last four lines in each table are for the sake of comparison of `FLAT` with `HORN` (see Section 6.3.2). R is for uniform and random landmark selection. K is for selecting the boundary vertices of a KAHIP partition as landmarks. We have used the version v0.71 of the KAHIP partitioning software [8], exploiting the `KaFFPa` algorithm, with the following parameters: The number of blocks to partition the graph was set to 178, so that we get (slightly more than) 2,000 landmarks. H is for a hybrid partition that initially creates a KAHIP partition (with half the landmarks) and then randomly chooses additional landmarks within each cell of the partition. IR indicates a variant of R that moves each randomly selected landmark to its closest important node. We have considered as "important" those nodes in the Berlin instance which are incident to road segments of category at most 3. SR indicates another variant of randomly selected landmarks, where each newly chosen random landmark excludes its closest 300 nodes (under the free-flow metric) from being landmarks in the future. As for the query algorithms, we used recursion budget 1 for `RQA` and we let $\texttt{FCA}^+$ settle the 6 closest landmarks, which is roughly the average number of settled landmarks by `RQA` as well.

6.3.2 `HORN` **@ Berlin.** Due to large space requirements, we could handle landmark hierarchies with up to 21,000 landmarks for the Berlin instance, which seems to be harder than that of Germany[1]. The average size per landmark in the hierarchy is 2.1MBytes. For a hierarchy of 10,256 landmarks the preprocessing of `HORN` took 5.0 hours, for 20,513 landmarks it took 9.7 hours, or at most 1.8sec per landmark in either case. The landmarks in each level of the hierarchy were chosen by the RANDOM (HR) and SPARSE-RANDOM (HSR) methods. We consider 4 levels of the hierarchy, according to the sizes of the landmarks' *areas of coverage*, i.e. the number of "nearby" destinations for which they possess travel-time summaries. The area of coverage for landmarks of level 4 is actually the entire graph. These are exactly the "global" landmarks which the corresponding variant of `FLAT` would also consider. The landmarks of the other levels have significantly smaller areas of coverage.

In summary, we created four distinct hierarchies, with 10,256 and 20,513 landmarks, based on HR and HSR landmark selection methods (cf. Table 7). The corresponding variants for `FLAT` possess the same (270 and 541, respectively) "global" landmarks. We created travel-time summaries using only the `TRAP` approximation method, within the subgraph induced by each landmark's area of coverage (extended, as in Assumption 2.3).

The experimental results of `FCA` for R_{270}, SR_{270}, R_{541} and SR_{541} are shown in Tables 5 and 6. Table 8 summarises the experimental evaluation of `HQA`. Interestingly,its performance with HR-landmarks is better than that with HSR-landmarks, probably because the ESC criterion seems more effective in the former case. Table 9 summarizes the comparison of `HORN` with `FLAT`. The results for HR_{10256}, HSR_{10256}, HR_{20513} and HSR_{20513} are compared with the corresponding results of `FCA` for R_{270}, SR_{270}, R_{541} and SR_{541} (i.e., with the same number of global landmarks), respectively. There is a significant improvement in query performance (e.g., more than 40% w.r.t. Dijkstra-ranks), but also in quality of the produced solution (by more than 41%), at the cost of increasing the space requirements by a factor of 6.44 at most.

6.3.3 `FLAT` **@ Germany.** We have tried `FLAT` with R-, K-, SR- and SK-landmark sets. Using 6 threads for the preprocessing phase, the average preprocessing time is less than 90sec and the average space is up to 25.7Mbytes, per landmark. We only used 2,000 landmarks which, as a fraction of the vertex set, is significantly smaller (by an order of magnitude) than in the instance of Berlin. Notably, for the K-landmark set `FLAT`'s query performances are now much worse than those for R-, SR- and SK-landmark sets, by almost an order of magnitude. This is probably due to the fact that the number of cells is now too small (36) so as to constrain ourselves to roughly 2,000 landmarks. On the other hand, for the R- SR-, and SK-landmark sets the query performances are quite remarkable, comparable to those for Berlin (by means of speedup), despite the significantly smaller density of the

[1] We observed that the speedups are significantly better in Germany, despite the fact that we consider the same number of landmarks in a larger, by an order of magnitude, network. This is probably due to stronger correlation of time-dependence among different road segments in an urban environment, rather than in a nationwide road network.

landmarks in the network. Tables 10 and 11 demonstrate the performance of the query algorithms with respect to absolute times and Dijkstra-ranks, respectively. E.g., the best speedup against TDD is achieved by SK-landmarks, and is more than $1,531$ in Dijkstra-ranks, and more than 902 in absolute query-times, with worst-case error at most 1.534%.

6.4 Comparison with [9]. We improve significantly our previous experimental evaluation of FLAT in the following sense. We significantly reduce the required preprocessing space, from 70GB down to at most 27GB for the instance of Berlin. This improvement allowed us to also conduct experiments for instance of Germany, which in this new implementation requires space of 51GB for the R-landmarks of size $2,000$. We provide novel landmark sets, which significantly improve the speedup (cf. Table 9). It is mentioned at this point that, since we have executed our new experiments on a different machine, we re-executed the experiments for R_{2000} and K_{2000}. This led to slightly worse running times, compared to the reported values in [9], which are nevertheless more accurate here since we indeed account for the exact execution times, whereas in our previous work we had to exclude the disk-IO accesses due to the limited memory-capabilities that we had at our disposal. Moreover, the relative errors and the Dijkstra-rank values in the present paper are better than those reported in [9], due to several improvements in our code. The best speedup factor is now achieved by SR-landmarks and FCA: The average query time is 81μsec, with worst-case relative error 0.771% and speedup more than $1,227$. With respect to the SK-landmarks, FCA achieves 83μsec, with worst-case relative error 0.781% and speedup more than $1,570$.

We also conducted extensive experimental evaluation for random queries of specific ranges, which confirmed our intuition that our query algorithms perform much better for long-range queries rather than for medium-range or short-range queries. This fact indeed motivated our construction of the HORN oracle. The evaluation of HORN for the instance of Berlin demonstrates that we can achieve remarkable speedups also for short-range and medium-range queries.

Finally, the creation of the temporal travel-time summaries as a response to live-traffic reporting is indeed done very fast, e.g., in less than 21sec for disruptions that would last about 15min.

6.5 Comparison with [4]. Our main goal in this work is to demonstrate the practicality of FLAT and HORN, which provide *provable* guarantees w.r.t. query times, stretch factors and preprocessing requirements, for large-scale real data sets. The strong aspects of our oracles are the simplicity of the query algorithms, the remarkably small (optimal in most cases) observed stretches, and the achieved speedups. E.g., for Germany FCA responds in 1.269msec, achieving a speedup against TDD more than 902 in absolute query-times and $1,531$ w.r.t. the *machine-independent* Dijkstra-ranks, and worst-case error less than 1.534%. This performance is clearly very encouraging and quite competitive against the current state-of-art approach of *time-dependent Contraction Hierarchies* [4]. On the negative side, the preprocessing space and time requirements still remains large, especially when compared to the small space requirements of the TCH-approach of [4].

7 Conclusions and Future Work

We provided an extensive experimental evaluation of landmark-based oracles for time-dependent road networks. We are currently exploring different landmark sets that will achieve even better speedups and/or approximation guarantees. As a first step, inspired by the *sparsification* in the R-landmark set, we also explored techniques for sparsifying also the KAHIP landmark sets. Preliminary results for Berlin have demonstrated remarkable results. E.g., we achieved average response times of 83μsec, worst-case error of 0.781% and Dijkstra-rank speedup more than $1,570$ for the Berlin instance.

In case that space is a main concern, we observed the full scalability in the treadoffs of our oracles between space and query-responses. E.g., consuming space 3.2GB we can achieve query-response times 0.73msec, relative error 2.198%, for the Berlin instance. We are currently experimenting with more sophisticated landmark schemes to further improve the space vs. query-responses tradeoff.

We are finally exploring further improvements in the compression schemes, and the exploitation of parallelism, not only as a simple load-balancing scheme, but also algorithmically, which will further reduce the requirements for preprocessing space and time, and live-traffic updating.

References

[1] R. Agarwal and P. Godfrey. Distance oracles for stretch less than 2. In *24th Symp. on Discrete Algorithms (SODA'13)*, pp. 526–538. ACM-SIAM, 2013.

[2] H. Bast, D. Delling, A. Goldberg, M. Müller-Hannemann, T. Pajor, P. Sanders, D. Wagner, and R. Werneck. Route planning in transportation networks. *CoRR*, abs/1504.05140, 2015.

[3] G. V. Batz, D. Delling, P. Sanders, and C. Vetter. Time-dependent contraction hierarchies. In *11th Workshop on Algorithm Engineering and Experiments (ALENEX'09)*, pp. 97–105, SIAM, April 2009.

[4] G. V. Batz, R. Geisberger, P. Sanders, and C. Vetter. Minimum time-dependent travel times with contraction hierarchies. *ACM Journal of Experimental Algorithmics*, 18(1.4):1–43, April 2013.

[5] D. Delling. Time-Dependent SHARC-Routing. *Algorithmica*, 60(1):60–94, May 2011.

[6] D. Delling and G. Nannicini. Core routing on dynamic time-dependent road networks. *Informs Journal on Computing*, 24(2):187–201, 2012.

[7] S. E. Dreyfus. An appraisal of some shortest-path algorithms. *Operations Research*, 17(3):395–412, 1969.

[8] KaHIP – Karlsruhe High Quality Partitioning, May 2014.

[9] S. Kontogiannis, G. Michalopoulos, G. Papastavrou, A. Paraskevopoulos, D. Wagner, and C. Zaroliagis. Analysis and experimental evaluation of time-dependent distance oracles. In *17th Workshop on Algorithm Engineering and Experiments (ALENEX 2015)*, pp. 147–158, 2015.

[10] S. Kontogiannis, D. Wagner, and C. Zaroliagis. Hierarchical oracles for time-dependent networks. *CoRR*, abs/1502.05222, 2015.

[11] S. Kontogiannis and C. Zaroliagis. Distance oracles for time-dependent networks. *Algorithmica*, 2015.

[12] G. Mali, P. Michail, A. Paraskevopoulos, and C. Zaroliagis. A New Dynamic Graph Structure for Large-Scale Transportation Networks. In *Algorithms and Complexity* (CIAC 2013), LNCS 7878:312–323, Springer,2013.

[13] METIS – serial graph partitioning and fill-reducing matrix ordering, 2013. Stable version: 5.1.0.

[14] G. Nannicini, D. Delling, L. Liberti, and D. Schultes. Bidirectional A* search on time-dependent road networks. *Networks*, 59:240–251, 2012.

[15] A. Orda and R. Rom. Shortest-path and minimum delay algorithms in networks with time-dependent edge-length. *Journal of the ACM*, 37(3):607–625, 1990.

[16] M. Patrascu and L. Roditty. Distance oracles beyond the Thorup-Zwick bound. *SIAM J. Comput.*, 43(1):300–311, 2014.

[17] E. Porat and L. Roditty. Preprocess, set, query! *Algorithmica*, 67(4):516–528, 2013.

[18] PTV AG – Planung Transport Verkehr. http://www.ptv.de.

[19] TomTom. http://www.tomtom.com/

[20] C. Sommer. Shortest-path queries in static networks. *ACM Computing Surveys*, 46, 2014.

[21] C. Sommer, E. Verbin, and W. Yu. Distance oracles for sparse graphs. In *50th Symp. on Foundations of Computer Science (FOCS'09)*, pages 703–712, IEEE 2009.

[22] M. Thorup and U. Zwick. Approximate distance oracles. *Journal of the ACM*, 52(1):1–24, 2005.

[23] C. Wulff-Nilsen. Approximate distance oracles with improved preprocessing time. In *23rd Symp. on Discrete Algorithms (SODA'12)*, pp. 202–208, ACM–SIAM, 2012.

[24] C. Wulff-Nilsen. Approximate distance oracles with improved query time. In *24th Symp. on Discrete Algorithms (SODA'13)*, pp. 539–549, ACM–SIAM 2013.

List of Tables

Symbol	Description
$[k]$	The set of integers $\{1, 2, \ldots, k\}$.
$G = (V, A)$	The graph representing the underlying road network. $n = \|V\|$ and $m = \|A\|$.
$diam(G, D)$	The diameter of G under an arc-cost metric D.
$\mathcal{P}_{o,d}$	Set of *od*-paths in G.
$p \bullet q$	The concatenation of the ux-path p with the xv-path q at vertex x.
$ASP[o, d](t_o)$	Set of $(1+\varepsilon)$-approximations of minimum-travel-time *od*-paths in G, for given departure-time $t_o \geq 0$.
$SP[o, d](t_o)$	Set of minimum-travel-time *od*-paths in G, for given departure-time $t_o \geq 0$.
$B[v](t_v)$	A ball growing from $(v, t_v) \in V \times [0, T)$, in the time-dependent metric, until either the destination d is reached or the closest landmark ℓ_v from (v, t_v) is settled.
$B[v; F](t_v)$	A ball growing from $(v, t_v) \in V \times [0, T)$, in the time-dependent metric, of size $F \in \mathbb{N}$.
$\overline{B}[v; F]$ / $\underline{B}[v; F]$	A ball growing from $v \in V$, in the full-congestion / free flow metric, of (integer) size $F \in \mathbb{N}$.
$\overline{B}[v; R]$ / $\underline{B}[v; R]$	A ball growing from $v \in V$, in the full-congestion / free flow metric, of (scalar) radius $R > 0$.
$B'[v; F](t_v)$	A ball growing from $(v, t_v) \in V \times [0, T)$, in the time-dependent metric, of size F polylog(F), according to Assumption 2.3.
$d[a](t)$	The limited-window arc-travel-time function for arc $a \in A$, with departure-time $t \in [0, T)$ for some *constant* time-period $T > 0$ (e.g., a single day).
M_a	Maximum possible travel-time ever seen at arc a.
M	Maximum arc-travel-time ever seen in any arc.
$D[a](t)$	Periodic arc-travel-time function for arc $a \in A$, with domain $t \in [0, \infty)$.
$Arr[a](t)$	The arc-arrival-time function for arc $a \in A$.
$D[o, d]$	Minimum-travel-time *function*, from o to d.
$\Gamma[o, d]$	Dijkstra-ranks *function*, from o to d.
$D_{\max}[o, d]$ / $D_{\min}[o, d]$	The maximum and minimum value of $D[o, d]$.
$\overline{D}[a]$ / $\underline{D}[a]$	Travel-times of a in *full-congestion* and *free-flow* metrics, respectively.
$\overline{\Delta}[o, d]$ / $\underline{\Delta}[o, d]$	An upper-approximating / lower-approximating function to $D[o, d]$.
$Arr[o, d]$	Earliest-arrival-time *function*, from o to d.
$t_u \quad (t_v)$	Departure-time from the tail u (arrival-time at the head v) for the arc $uv \in A$.
$TDSP(o, d, t_o)$	The problem of finding an min-cost *od*-path, given a departure-time t_o.
$TDSP(o, \star, t_o)$	The problem of finding a min-cost paths tree from o, given a departure-time t_o.
$TDSP(o, d)$	The problem of constructing a succinct representation of min-cost *od*-paths *function*.
K_a	Number of breakpoints in the arc-travel-time function $D[a]$.
K	Total number of breakpoints in the arc-travel-time functions.
$K_{\max}$	The maximum number of breakpoints, among the arc-travel-time functions.
K^*	Total number of *concavity-spoiling* breakpoints in the arc-travel-time functions.
$\Lambda_{\max}$	Maximum slope among minimum-travel-time functions.
$\Lambda_{\min}$	Absolute value of minimum slope among minimum-travel-time functions.
ζ	ratio of min-travel-times in opposite directions between two vertices for any specific departure-time.
r	The recursion budget for `RQA` and `HQA`.
`BIS`	The bisection approximation method for minimum-travel-time functions.
`TRAP`	The trapezoidal approximation method for minimum-travel-time functions.
`FCA`	The *Forward Constant Approximation* query algorithm.
`RQA`	The *Recursive Query Algorithm*.
`HQA`	The *Hierarchical Query Algorithm*, based on a hierarchy of landmarks.
`FLAT`	The oracle that uses landmarks possessing summaries towards all possible destinations.
`HORN`	The oracle that uses a hierarchy of landmarks and the `HQA` query algorithm.

Table 1: Summary of notation.

PARAMETER \ INSTANCE	Berlin (TomTom)	Germany (PTV AG)
#Nodes	473,253	4,692,091
#Edges	1,126,468	11,183,060
Time Period	24h (Tue)	24h (Tue-Wed-Thu)
$\lambda_{\max}$	0.017	0.130
$-\lambda_{\min}$	-0.013	-0.130
#Arcs with constant traversal-times	924,254	10,310,234
#Arcs with non-constant traversal-times	20,2214	872,826
Min #Breakpoints	4	5
Avg #Breakpoints	10.4	16.3
Max #Breakpoints	125	52
Total #Breakpoints	3,234,213	25,424,506

Table 2: Arc-travel-time function statistics. $\lambda_{\max}$ and $-\lambda_{\min}$ are the maximium and minimum slope among the arc-traversal-time functions.

PARAMETER \ INSTANCE	Berlin (TomTom)	Germany (PTV AG)
ζ_{avg}	1.008	1.004
$\zeta_{\max}$	1.189	1.050
$\Lambda_{\max}$	0.190	0.216
$-\Lambda_{\min}$	-0.150	-0.197
Max Offset	46,039	51,850
Min Offset	-1,472	-6,331
Max Min-Travel-Time	43,797	48,382

Table 3: Min-cost path-travel-time statistics.

	ExtFF Rank @ Berlin			*ExtFF Rank* @ Germany		
FF Rank	avg	max	blow-up (%)	avg	max	blow-up (%)
100	136.9	671	6.710	135.9	670	6.700
1,000	1,532.7	3,898	3.898	1,661.7	8,299	8.299
4,000	6,315.0	15,552	3.888	7,110.2	16,024	4.006
6,000	9,669.1	25,951	4.326	10,783.3	29,145	4.858
10,000	16,142.8	36,744	3.675	17,890.1	42,415	4.242

Table 4: Free-flow blow-up statistics for the instances of Berlin and Germany. The worst-case blow-up factor is less than 8.3 times.

	TDD		FCA		FCA$^+$(6)		RQA	
	Time (msec)	Rel.Error %	Time (msec)	Rel.Error %	Time (msec)	Rel.Error %	Time (msec)	Rel.Error %
R_{2000}	92.906	0	0.100	0.969	0.527	0.405	0.519	0.679
K_{2000}			0.115	1.089	0.321	0.405	0.376	0.523
H_{2000}			0.102	0.886	0.523	0.332	0.445	0.602
IR_{2000}			0.086	0.923	0.489	0.379	0.473	0.604
SR_{2000}			0.081	0.771	0.586	0.317	0.443	0.611
SK_{2000}			0.083	0.781	0.616	0.227	0.397	0.464
R_{541}			0.326	1.854	1.887	0.693	1.904	1.610
SR_{541}			0.451	1.638	3.252	0.614	2.856	1.531
R_{270}			0.639	2.583	3.707	0.881	3.842	2.482
SR_{270}			0.730	2.198	4.491	0.745	4.271	2.336

Table 5: Performance of FCA, FCA$^+$(6) and RQA, w.r.t. the *running times* and *relative errors*, at 2.64sec resolution, for a query set of 10,000 *random queries* in Berlin.

	TDD		FCA		FCA$^+$(6)		RQA	
	Rank	Speedup	Rank	Speedup	Rank	Speedup	Rank	Speedup
R_{2000}	146,022	1	150	973.480	877	166.502	925	157.862
K_{2000}			190	768.537	866	168.616	670	217.943
H_{2000}			154	948.195	851	171.589	777	187.931
IR_{2000}			135	1,081.644	823	177.426	839	174.043
SR_{2000}			119	1,227.075	952	153.384	776	188.173
SK_{2000}			93	1,570.129	755	193.406	501	291.461
R_{541}			545	267.930	3,178	45.947	3,406	42.872
SR_{541}			638	228.874	3,684	39.637	3,950	36.967
R_{270}			1,075	135.834	6,198	23.559	6,702	21.788
SR_{270}			1,195	122.194	7,362	19.835	7,398	19.738

Table 6: Performance of FCA, FCA$^+$(6) and RQA, w.r.t. *Dijkstra ranks*, at 2.64sec resolution, for a query set of 10,000 *random queries* in Berlin.

Level	Size of Levels		Area of coverage	Excluded Ball Size (for HSR)	
	$\|L\| = 10,256$	$\|L\| = 20,513$		$\|L\| = 10,256$	$\|L\| = 20,513$
L_1	7,685	15,370	1,274	35	15
L_2	1,604	3,208	29,243	150	80
L_3	697	1,394	154,847	350	180
L_4	270	541	292,356	800	400

Table 7: Landmark hierarchies for HORN, based on HR and HSR landmark selection methods, for the Berlin instance.

	TDD				HQA			
	Time (msec)	Rel.Error %	Rank	Speedup	Time (msec)	Rel.Error %	Rank	Speedup
HR_{10256}	92.906	0	146,022	1	0.354	1.499	636	229.594
HSR_{10256}					0.436	1.409	721	202.527
HR_{20513}					0.217	1.051	324	450.685
HSR_{20513}					0.314	0.919	378	386.302

Table 8: Performance of HQA, w.r.t. the *running times*, *relative errors* and *Dijkstra ranks*, at 2.64sec resolution, for a query set of 10,000 *random queries* in Berlin.

	Improvement in			Deterioration in
	Query Times (%)	Worst-case Relative Error (%)	Dijkstra Ranks (%)	Space (times)
R_{270} vs HR_{10256}	44.60	41.96	40.83	6.089
SR_{270} vs HSR_{10256}	40.27	35.89	39.66	6.407
R_{541} vs HR_{20513}	33.43	43.31	40.55	6.195
SR_{541} vs HSR_{20513}	30.37	43.89	40.75	6.438

Table 9: Comparison of HQA versus FCA in Berlin. Two global-landmark sizes (270 and 541) and two landmark-selection methods (RANDOM and SPARSE-RANDOM) are considered.

	TDD		FCA		FCA$^+$(6)		RQA	
	Time (msec)	Rel.Error %	Time (msec)	Rel.Error %	Time (msec)	Rel.Error %	Time (msec)	Rel.Error %
R_{2000}	1,145.060	0	1.532	1.567	8.529	0.742	9.219	1.502
K_{2000}			10.455	2.515	15.209	1.708	30.577	2.343
SR_{2000}			1.275	1.444	9.952	0.662	9.011	1.412
SK_{2000}			1.269	1.534	9.689	0.676	7.653	1.475

Table 10: Performance of FCA, FCA$^+$(6) and RQA, w.r.t. the *running times* and *relative errors*, at 17.64sec resolution, for a query set of 10,000 *random queries* in Germany.

	TDD		FCA		FCA$^+$(6)		RQA	
	Rank	Speedup	Rank	Speedup	Rank	Speedup	Rank	Speedup
R_{2000}	1,717,793	1	1,659	1,035.439	10,159	169.091	11,045	155.527
K_{2000}			9,302	184.669	15,373	111.741	30,137	56.999
SR_{2000}			1,277	1,345.178	9,943	172.764	9,182	187.082
SK_{2000}			1,122	1,531.010	9,000	190.866	7,975	215.397

Table 11: Query performance of FCA, FCA$^+$(6) and RQA, w.r.t. Dijkstra ranks, at 17.64sec resolution, for a query set of 10,000 *random queries* in Germany.